arXiv:1410.6556v1 [stat.ME] 24 Oct 2014

Forward variable selection for sparse ultra-high dimensional varying coefficient models

Ming-Yen Cheng, Toshio Honda, and Jin-Ting Zhang

Abstract

Varying coefficient models have numerous applications in a wide scope of scientific areas. While enjoying nice interpretability, they also allow flexibility in modeling dynamic impacts of the covariates. But, in the new era of big data, it is challenging to select the relevant variables when there are a large number of candidates. Recently several work are focused on this important problem based on sparsity assumptions; they are subject to some limitations, however. We introduce an appealing forward variable selection procedure. It selects important variables sequentially according to a sum of squares criterion, and it employs an EBIC- or BIC-based stopping rule. Clearly it is simple to implement and fast to compute, and it possesses many other desirable properties from both theoretical and numerical viewpoints. We establish rigorous selection consistency results when either EBIC or BIC is used as the stopping criterion, under some mild regularity conditions. Notably, unlike existing methods, an extra screening step is not required to ensure selection consistency. Even if the regularity conditions fail to hold, our procedure is still useful as an effective screening procedure in a less restrictive setup. We carried out simulation and empirical studies to show the efficacy and usefulness of our procedure.

Keywords: B-spline; EBIC; independence screening; marginal model; semi-varying coefficient models; sub-Gaussion error; structure identification.

Ming-Yen Cheng is Professor, Department of Mathematics, National Taiwan University, Taipei 106, Taiwan (Email: cheng@math.ntu.edu.tw). Toshio Honda is Professor, Graduate School of Economics, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo 186-8601, Japan (Email: t.honda@r.hit-u.ac.jp). Jin-Ting Zhang is Associate Professor, Department of Statistics & Applied Probability, National University of Singapore, 3 Science Drive 2, Singapore 117546 (Email: stazjt@nus.edu.sg). This research is partially supported by the Hitotsubashi International Fellow Program and the Mathematics Division, National Center of Theoretical Sciences (Taipei Office). Cheng is supported by the Ministry of Science and Technology grant MOST101-2118-M-002-001-MY3. Honda is supported by the JSPS Grant-in-Aids for Scientific Research (A) 24243031 and (C) 25400197. Zhang is supported by the National University of Singapore research grant R-155-000-128-112.

1 Introduction

We consider variable selection problem for the varying coefficient model defined by

$$Y = \sum_{j=0}^{p} \beta_{0j}(T) X_j + \epsilon, \tag{1}$$

where Y is a scalar response variable, $X_0 \equiv 1$, $X_1, \ldots, X_p$ are the candidate covariates, ϵ is the random error, and $T \in [0, 1]$. The coefficient functions β_{0j}, $j = 0, 1, \ldots, p$, are assumed to vary smoothly with T, and are non-zero for only a subset of the p candidate covariates. The variable T is an influential variable, such as age or income in econometric studies, and is sometimes called the index variable. The varying coefficient model is a popular and useful semiparametric approach to modeling data that may not obey the restrictive form of traditional parametric models. In particular, while it retains the nice interpretability of the linear models, it allows good flexibility in capturing the dynamic impacts of the relevant covariates on the response Y. In addition, in practical applications, some of the true covariates may have simply constant effects while the others have varying effects. Such situations can be easily accommodated by a variant, the so called semi-varying coefficient model [31, 34]. Furthermore, model (1) has been generalized to modeling various data types including count data, binary response, clustered/longitudinal data, time series, and so on. We refer to [13] for a comprehensive review and the extensive literature.

Due to recent rapid developments in technology for data acquisition and storage, nowadays a lot of high-dimensional data sets are collected in various research fields where varying coefficient models find meanings and applications, such as medicine, marketing and so on. In such situations, the model used to analyze the data is usually sparse, that is, the number of true covariates is not large even when the dimension is very large. Therefore, under the sparsity condition, some effective variable selection procedures are necessary in order to carry out meaningful statistical estimation and inference. In this regard, the penalized variable selection approach emerged as the mainstream in the recent decade. Existing general penalty functions for sparse (ultra-)high-dimensional models include the Lasso [27], group Lasso [21, 32], adaptive Lasso [36], SCAD [8] and Dantzig selector [3].

In ultra-high dimensional cases where the dimensionality p is very large, selection consistency becomes challenging and nearly impossible for existing variable selection methods to achieve, however. Thus, an additional independence screening step is usually necessary before variable selection is carried out. For example, sure independence

screening (SIS) methods are introduced by [9] and [11] for linear models and generalized linear models respectively, and nonparametric independence screening (NIS) is suggested for additive models by [7]. Under general parametric models, [12] suggested using the Lasso at the screening stage before implementing a local linear approximation to the SCAD (or general folded concave) penalty at the second stage. In all of the above mentioned variable selection and independence screening methods, some tuning parameter or threshold value is involved which needs to be determined by the user or by some elaborated means. Under the considered varying coefficient model (1), there are some existing work on penalized variable selection in several different setups of the dimensionality p, using the Lasso or folded concave penalties such as the SCAD [1, 17, 22, 26, 28, 29, 30]. In ultra-high dimensional cases, for the independence screening purpose, the Lasso is recommended by [29] and NIS is considered by several authors [5, 10, 19, 25]. Again, all of these methods require selection of some tuning parameter or threshold value.

More recently, an alternative forward variable selection approach receives increasing attention for linear regression. The literature along this line includes the least angle regression (LAR) [6], the forward iterative regression and shrinkage technique (FIRST) [16], the forward Lasso adaptive shrinkage (FLASH) [23], and the sequential Lasso (SLASSO) [20]. Such methods enjoy desirable theoretical properties, including selection consistency, and have advantages from numerical aspects. Motivated by the above observations, we propose and investigate thoroughly a forward variable selection procedure for the considered varying coefficient model in ultra-high dimensional covariate cases, where the dimensionality can be much larger than the sample size. The proposed method is constructed in a spirit similar to the SLASSO [20], which employes Lasso in the forward selection and uses the EBIC [4] as the stopping criterion. However, the selection criterion of our method is based on the reduction in the sum of squared residuals, instead of the Lasso. This is because our preliminary simulation studies suggested that the proposed one performs better than the analogue of the Lasso for the varying coefficient model considered here.

The stopping rule of the proposed forward selection procedure is based on the analogue of the EBIC [4], or alternatively the BIC, for the varying coefficient model. The consistency result of the EBIC for model selection in ultra-high dimensional additive models is established by [18] when the number of true covariates p_0 is bounded. The paper also assumes some knowledge of the number of true covariates, which may be unrealistic or difficult to obtain in some cases. On the other hand, without this kind of

knowledge, the number of all possible subsets of the candidate variables to be considered is too large and there is no guarantee that EBIC-based model selection will perform properly. Therefore, it makes sense to consider a forward selection procedure, which does not require such prior knowledge, and use the EBIC as the stopping criterion.

Suppose we have n i.i.d. observations $\{(\mathbf{X}_i, T_i, Y_i)\}_{i=1}^n$, where $\mathbf{X}_i = (X_{i0}, X_{i1}, \ldots, X_{ip})$, taken from the varying coefficient model (1):

$$Y_i = \sum_{j=0}^{p} \beta_{0j}(T_i) X_{ij} + \epsilon_i,\ i = 1, \ldots, n. \tag{2}$$

In our theoretical study, we deal with the ultra-high dimensional case where

$$\log p = O(n^{1-c_p}/L). \tag{3}$$

Here, c_p is a positive constant and L is the dimension of the B-spline basis used in the estimation of the coefficient functions. We will give more details on the B-spline basis and specify more conditions on p later in Sections 2 and 3; especially see Assumptions B(2) and B(3) for the conditions on p. Throughout this paper, $\#A$ denotes the number of elements of a set A, and A^c is the complement of A. We write S_0 for the set of indexes of the true covariates in model (1), that is, $\beta_{0j} \not\equiv 0$ for $j \in S_0$ and $\beta_{0j} \equiv 0$ for $j \in S_0^c$. In addition, we write p_0 for the number of true covariates, i.e. $p_0 \equiv \#S_0$, and consider the case that

$$p_0 = O((\log n)^{c_S}) \tag{4}$$

for some positive constant c_S. Here, condition (4) on p_0 is imposed for simplicity of presentation; it can be relaxed at the expense of restricting slightly the order of the dimension p specified in (3).

Under some assumptions we establish the selection consistency of our forward variable selection method when p can be larger than n and p_0 can grow slowly with n, as specified in (3) and (4). Importantly, this means that no independence screening is required before the proposed variable selection procedure. This nice property may be intuitively correct when dealing with sparse parametric models using methods like the SLASSO [20]. But it is not obvious for varying coefficient models; in model (1) each of the coefficient functions is modeled nonparametrically and involves L parameters in its spline estimation. We exploit desirable properties of B-spline bases to drive these strong theoretical results. Note also that our selection consistency results hold when either the EBIC or the BIC is used in the stopping rule.

Interestingly, contradictory to what is suggested for linear models, our simulation results indicate that for the considered varying coefficient model (1) the BIC outperforms the EBIC when they are used as the stopping criterion in the forward selection procedure. In fact, the EBIC stopping rule tends to stop the forward selection too early and make it miss some important variables. The reason behind this is that the penalty on adding another variable is too large. Some adjustments may be helpful in coping with this issue, but fortunately we can circumvent it by using simply the BIC and our simulation results show it works very well. Another problem worth of further study is whether the EBIC is really better in forward selection; it is to account for the large number of possible choices in model selection, but this issue vanishes in forward selection.

As mentioned earlier, there exist some useful procedures for variable selection in varying coefficient modeling. Nonetheless, the proposed method has many merits compared to them, from both practical and theoretical viewpoints. First, since the important variables are selected sequentially, the final model has good *interpretability* in the sense that we can rank the importance of the variables according to the order they are selected. Second, in practice we may have some *a priori* knowledge that certain relevant variables should be included in the model. In this case, we always have the *flexibility* to start from any subset that contains them. Third, our method employs reasonable sequential selection and stopping rules, and no tuning parameters or threshold parameters are present, meaning that the implementation and the computation are *simple and fast*. Fourth, there is a drastic gain in terms of *numeric stability* as no inversion of large matrices is necessary, as long as the number of true covariates p_0 is not large. By comparison, existing variable selection methods all require independence screening in advance, but the NIS and the group Lasso tend to choose many covariates in order not to miss any true covariates; thus inversion of large matrices is inevitable. (Notice that the spline estimation of each of the coefficient functions involves L number of parameters, which has to diverge to infinity with n, and we have only one observation for each subject in the present setup.) Fifth, same as [5], we improve on the order of p as compared with the conditions in [10]. In other words, the forward procedure can reduce the dimensionality more effectively. Finally, our method requires *milder regularity conditions* than the sparse Riesz condition [29] and the restricted eigenvalue conditions [2] for the Lasso, which are related to all the candidate covariates (Then, there may be a large set of "ill-behaved" covariates with indexes outside of S_0, especially when p is very large).

The assumptions we impose in Section 3 for the selection consistency of our method

may fail to hold in some cases. Nevertheless, in that case we can still use the proposed procedure for the purpose of independence screening, under a less restrictive setup specified in Section 2.4. Then, we will successfully reduce the number of covariates to a moderate order. This allows us to identify consistently the true covariates in the next stage, by applying the group SCAD or the adaptive group Lasso procedure to the variables that pass the screening. See Sections 2.4 and 3 for the details. Besides, some of the coefficient functions may be constant i.e. $\beta_{0j} \equiv \text{const}$ for some $j \in S_0$. Under such circumstances, we can carry out some group SCAD or adaptive Lasso procedures to detect both the constant coefficients and the varying coefficients, as suggested in Section 3 of [5]. We refer to [5] for such a two-stage approach, i.e. screening and then structure identification, and the theoretical and numerical justifications. Note that, there are indeed some advantages in using the proposed forward procedure as a screening tool. In particular, it tends to remove more irrelevant variables than NIS approaches do, and thus reducing the dimensionality more effectively. See Section 4.2 for some numerical comparisons.

This paper is organized as follows. In Section 2, we describe the proposed forward variable selection procedure. At each step, it uses the residual sum of squares resulted from spline estimation of an extended marginal model to determine the next candidate feature, and it uses the EBIC or the BIC to decide whether to stop or to include the newly selected feature and continue. We state the assumptions and theoretical results in Section 3. Results of simulation and empirical studies are presented in Section 4. Proofs of all the theoretical results are given in Section 5.

2 Method

In this section, we describe the proposed forward feature selection procedure. Before that, we introduce some notation. We write $\|f\|_{L_2}$ and $\|f\|_\infty$ for the L_2 and sup norm of a function f on $[0,1]$, respectively. When g is a function of some random variable(s), we define the L_2 norm of g by $\|g\| = [\mathrm{E}\{g^2\}]^{1/2}$. For a k-dimensional vector $\mathbf{x}$, $|\mathbf{x}|$ stands for the Euclidean norm and $\mathbf{x}^T$ is the transpose. We use the same symbol for transpose of matrices.

Recall S_0 is the set of true covariates in the varying coefficient model (1). Suppose that we have selected covariates sequentially and obtain index sets $S_1, \ldots, S_k$ as follows:

$$S_1 \subset S_2 \subset \cdots \subset S_k \equiv S \subset S_0.$$

That is, S_j is the index set of the selected covariates upon the completion of the jth step, for $j = 1, \ldots, k$. Note that S_1 can be the empty set ϕ, $\{0\}$ which corresponds to the intercept function, or some non-empty subset of S_0 given according to some *a priori* knowledge. Then, at the current $(k+1)$th step, we need to choose another candidate from S^c, and then we need to decide whether we should stop or add it to S and go to the next step. Our forward feature selection criterion is defined in (11), and we employ a version of the EBIC, given in (13), as the stopping rule. See [4] for more details about the EBIC.

2.1 Extended marginal model

In this section, we consider spline estimation of the extended marginal model when we add another index to the current index set S, which we will make use of in deriving our forward selection criterion. Hereafter we write $S(l)$ for $S \cup \{l\}$ for any $l \in S^c$. Temporarily we consider the following extended marginal model for $S(l), l \in S^c$:

$$Y = \sum_{j \in S(l)} \overline{\beta}_j(T) X_j + \epsilon_{S(l)}. \tag{5}$$

Here, the coefficient functions $\overline{\beta}_j,\ j \in S(l)$, are defined in terms of minimizing the following mean squared error with respect to $\beta_j,\ j \in S(l)$,

$$\mathrm{E}\Big\{ \Big(Y - \sum_{j \in S(l)} \beta_j(T) X_j \Big)^2 \Big\},$$

where the minimization is over the set of L_2 integrable functions on $[0,1]$. Note that $\|\overline{\beta}_j\|_{L_2}$ should be larger when $j \in S_0 - S$ than when $j \in S_0^c$. We will impose some assumptions on these coefficient functions later in this section and in Section 3.

First, we introduce some more notation related to the B-spline basis used in estimating the extended marginal model (5). Let $\mathbf{B}(t)$ denote the L-dimensional equi-spaced B-spline basis on $[0,1]$. We assume that $L = c_L n^{\kappa_L}$ where $\kappa_L \geq 1/5$. The order of the B-spline basis should be taken larger than or equal to two, under our smoothness assumptions on the coefficient functions in model (5). Assumptions B(4)-(5) given in Section 3 ensure that we can approximate the coefficient functions with the B-spline bases. See [24] for the definition of B-spline bases. We write

$$\begin{aligned} &\boldsymbol{W}_{ij} = \mathbf{B}(T_i) X_{ij} \in \mathbb{R}^L, \quad \boldsymbol{W}_{iS} = (\boldsymbol{W}_{ij}^T)_{j \in S}^T \in \mathbb{R}^{L \# S}, \\ &\boldsymbol{W}_j = (\boldsymbol{W}_{1j}, \ldots, \boldsymbol{W}_{nj})^T \quad \text{and} \quad \boldsymbol{W}_S = (\boldsymbol{W}_{1S}, \ldots, \boldsymbol{W}_{nS})^T. \end{aligned}$$

Note that $\boldsymbol{W}_{ij}$ is a vector of regressors in the spline estimation of $\overline{\beta}_j$ in model (5), and $\boldsymbol{W}_j$ and $\boldsymbol{W}_S$ are respectively $n \times L$ and $n \times (L\#S)$ matrices. Based on the B-spline basis, we can approximate the varying coefficient model (2) by the following approximate regression model:

$$Y_i = \sum_{j=0}^{p} \boldsymbol{\gamma}_{0j}^T \boldsymbol{W}_{ij} + \epsilon_i', \ i = 1, \ldots, n, \tag{6}$$

where $\boldsymbol{\gamma}_{0j} \in \mathbb{R}^L$ and $\boldsymbol{\gamma}_{0j}^T \mathbf{B}(t) \approx \beta_{0j}(t)$, $j = 0, 1, \ldots, p$. Similarly, the spline approximation model when the data come from the extended marginal model (5) is given by

$$Y_i = \sum_{j \in S(l)} \overline{\boldsymbol{\gamma}}_j^T \boldsymbol{W}_{ij} + \epsilon_{iS(l)}' = \overline{\boldsymbol{\gamma}}_S^T \boldsymbol{W}_{iS} + \overline{\boldsymbol{\gamma}}_l^T \boldsymbol{W}_{il} + \epsilon_{iS(l)}', \ i = 1, \ldots, n, \tag{7}$$

where $\overline{\boldsymbol{\gamma}}_S^T = (\overline{\boldsymbol{\gamma}}_j^T)_{j \in S}$ and $\overline{\boldsymbol{\gamma}}_j$, $j \in S(l)$, are defined by minimizing with respect to $\boldsymbol{\gamma}_j \in \mathbb{R}^L$, $j \in S(l)$, the following mean squared spline approximation error:

$$\mathrm{E}\Big\{ \sum_{i=1}^{n} \big(Y_i - \sum_{j \in S(l)} \boldsymbol{\gamma}_j^T \boldsymbol{W}_{ij}\big)^2 \Big\} = \mathrm{E}\Big\{ \big| \boldsymbol{Y} - \boldsymbol{W}_S \boldsymbol{\gamma}_S - \boldsymbol{W}_l \boldsymbol{\gamma}_l \big|^2 \Big\}$$

with $\boldsymbol{\gamma}_S^T = (\boldsymbol{\gamma}_j^T)_{j \in S}$. Note that $\overline{\boldsymbol{\gamma}}_j^T \mathbf{B}(t)$ should be close to the coefficient function $\overline{\beta}_j(t)$ in the extended marginal model (5). In particular, when $l \in S_0$, $\|\overline{\beta}_l\|_{L_2}$ should be large enough, and thus $|\overline{\boldsymbol{\gamma}}_l|$ should be also large enough.

We can estimate the vector parameters $\overline{\boldsymbol{\gamma}}_j$, $j \in S(l)$, in model (7) by the ordinary least squares estimates, denoted by $\widehat{\boldsymbol{\gamma}}_j$, $j \in S(l)$. Let $\widehat{\boldsymbol{W}}_{lS}$ and $\widehat{\boldsymbol{Y}}_S$ denote respectively the orthogonal projections of $\boldsymbol{W}_{lS}$ and $\boldsymbol{Y} = (Y_1, \ldots, Y_n)^T$ onto the linear space spanned by the columns of $\boldsymbol{W}_S$, that is,

$$\widehat{\boldsymbol{W}}_{lS} = \boldsymbol{W}_S (\boldsymbol{W}_S^T \boldsymbol{W}_S)^{-1} \boldsymbol{W}_S^T \boldsymbol{W}_l \quad \text{and} \quad \widehat{\boldsymbol{Y}}_S = \boldsymbol{W}_S (\boldsymbol{W}_S^T \boldsymbol{W}_S)^{-1} \boldsymbol{W}_S^T \boldsymbol{Y} \, .$$

Note that $\widehat{\boldsymbol{W}}_{jS}$ is an $n \times L$ matrix. Then the ordinary least square estimate of $\overline{\boldsymbol{\gamma}}_l$, denoted by $\widehat{\boldsymbol{\gamma}}_l$, can be expressed as

$$\widehat{\boldsymbol{\gamma}}_l = (\widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{W}}_{lS})^{-1} \widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{Y}}_S, \tag{8}$$

where $\widetilde{\boldsymbol{W}}_{jS} = \boldsymbol{W}_j - \widehat{\boldsymbol{W}}_{jS}$ and $\widetilde{\boldsymbol{Y}}_S = \boldsymbol{Y} - \widehat{\boldsymbol{Y}}_S$. Note that $\widehat{\boldsymbol{\gamma}}_l^T \mathbf{B}(t)$ is the spline estimate of the coefficient function $\overline{\beta}_l(t)$ in the extended marginal model (5).

2.2 Forward feature selection procedure

Recall that at the current step we are given S, the index set of the covariates already selected, and the job is to choose from S^c another candidate and then decide whether

we should add it to S or we should not and stop. For the purpose of forward feature selection, we consider the reduction in the sum of squared residuals, or equivalently the difference in the variance estimation, when adding l to S. Specifically, we compute $\widehat{\sigma}_S^2 - \widehat{\sigma}_{S(l)}^2$, where $\widehat{\sigma}_Q^2$ is the variance estimate for a subset of covariates indexed by Q given as

$$\widehat{\sigma}_Q^2 = \frac{1}{n}\Big\{\boldsymbol{Y}^T\boldsymbol{Y} - \boldsymbol{Y}^T\boldsymbol{W}_Q(\boldsymbol{W}_Q^T\boldsymbol{W}_Q)^{-1}\boldsymbol{W}_Q^T\boldsymbol{Y}\Big\}. \tag{9}$$

Using (8), we can rewrite $\widehat{\sigma}_S^2 - \widehat{\sigma}_{S(l)}^2$ as

$$\begin{aligned}\widehat{\sigma}_S^2 - \widehat{\sigma}_{S(l)}^2 &= \frac{1}{n}\big(\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{Y}}_S\big)^T\big(\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}\big)^{-1}\big(\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{Y}}_S\big) \\ &= \widehat{\boldsymbol{\gamma}}_l^T\Big(\frac{1}{n}\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}\Big)\widehat{\boldsymbol{\gamma}}_l \approx \mathrm{E}\Big\{\big(\overline{\beta}_l(T)\widetilde{X}_{lS}\big)^2\Big\},\end{aligned} \tag{10}$$

where $\widetilde{X}_{lS} = X_l - \widehat{X}_{lS}$ and $\widehat{X}_{lS}$ is the projection of X_l to $\big\{\sum_{j\in S}\beta_j(T)X_j\big\}$ with respect to the L_2 norm $\|\cdot\|_{L_2}$.

As noted earlier, if $l \in S_0$ then $\|\overline{\beta}_l\|_{L_2}$ will be large enough. Furthermore, $n^{-1}\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}$ will have desirable properties under Assumption X(2) given in Section 3; see Lemma 1 for the details. Hence, following from expression (10) and recalling that $\widehat{\boldsymbol{\gamma}}_l^T\mathbf{B}(t)$ is the spline estimate of $\overline{\beta}_l(t)$, we choose the candidate index as

$$l^* = \underset{l\in S^c}{\operatorname{argmin}}\,\widehat{\sigma}_{S(l)}^2\,. \tag{11}$$

Then, we have high confidence that l^* belongs to $S_0 - S$ provided that the latter is non-empty, and we take X_{l^*} as the next candidate feature. At first, instead of (11), we considered choosing

$$l^\dagger = \underset{j\in S^c}{\operatorname{argmax}}\,\big|\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{Y}}_S\big| \tag{12}$$

as the next candidate index, as motivated by the sequential Lasso for linear models proposed by [20]. However, after some simulation studies we found that, contrary to the nice properties of its counterpart in linear models, (11) performs better for the varying coefficient model we study.

To determine whether or not to include the candidate feature X_{l^*} in the set of selected ones, we employ the EBIC criterion. Specifically, we define the EBIC of a subset of covariates indexed by Q as the following:

$$\mathrm{EBIC}(Q) = n\log(\widehat{\sigma}_Q^2) + \#Q \times L(\log n + 2\eta\log p), \tag{13}$$

where η is a fixed constant and $\widehat{\sigma}_Q^2$ is given in (9). Then, at the current $(k+1)$th step, we should select the new covariate X_{l^*} with l^* defined in (11), provided that the EBIC

decreases when we add l^* to S and form $S(l^*)$. Otherwise, if the EBIC increases, we should not select any more covariates and stop at the kth step. Note that the EBIC defined in (13) reduces to the BIC when η is taken as 0. And, the theoretical results given in Section 3, in particular the consistency results given in Theorem 2, hold when either the EBIC or the BIC is used as the stopping criterion in the proposed method. In the following, we define formally the proposed forward feature selection algorithm.

Forward feature selection algorithm.

Initial step: Specify S_1, which can be taken as the empty set ϕ, $\{0\}$, or some non-empty subset of S_0 chosen based on some *a priori* knowledge, and compute EBIC(S_1).

Sequential selection: At the $(k+1)$th step, compute $\widehat{\sigma}^2_{S_k(l)}$ for every $l \in S_k^c$, and find

$$l^*_{k+1} = \underset{l \in S_k^c}{\operatorname{argmin}} \widehat{\sigma}^2_{S_k(l)} \, .$$

Then, let $S_{k+1} = S_k \cup \{l^*_{k+1}\}$ and compute EBIC(S_{k+1}). Stop and declare S_k as the set of selected covariate indexes if EBIC(S_{k+1}) $>$ EBIC(S_k); otherwise, change k to $k+1$ and continue to search for the next candidate feature.

The forward procedure with the EBIC stopping rule tends to stop a little too early and miss some relevant variables, and we need some kind of modification when we implement it. For example, some adjustment of the degrees of freedom will be helpful. All the details are given in Section 4.

2.3 Sparsity assumptions

We need some assumptions to establish consistency of the proposed procedure, especially Assumption B(1) given below. When conditions B(1)-(2) are not fulfilled, another setup in which we can use the proposed method as a screening approach is given in Section 2.4. In this paper, C_1, C_2, ... are generic positive constants and their values may change from line to line. Recall that S_0 is the index set of the true variables in model (1).

Assumption B(1)-(2)

B(1) For some large positive constant C_{B1},

$$\max_{j \in S_0 - S} \|\overline{\beta}_j\|_{L_2} / \max_{j \in S_0^c} \|\overline{\beta}_j\|_{L_2} > C_{B1}$$

uniformly in $S \subsetneq S_0$. Note that C_{B1} should depend on the other assumptions on the covariates, specifically Assumptions **X** and **T** given in Section 3.

B(2) Set $\kappa_n = \inf_{S \subsetneq S_0} \max_{j \in S_0 - S} \|\overline{\beta}_j\|_{L_2}$. We assume

$$\frac{n\kappa_n^2}{L \max\{\log p, \log n\}} > n^{c_\beta} \quad \text{and} \quad \kappa_n > \frac{L}{n^{1-c_\beta}}$$

for some small positive constant c_β. In addition, if $\eta = 0$ in (13) i.e. if BIC is used, we require that

$$\frac{L \log n}{\log p} \to \infty \,.$$

An assumption similar to Assumption B(1) is imposed in [20] and such assumptions are inevitable in establishing the selection consistency of forward procedures. These assumptions ensure that the chosen index l^*, given in (11), will be from $S_0 - S$. When such assumptions fail to hold, our method may choose some covariates from S_0^c. However, these covariates will be removed at the second stage mentioned in the Introduction. See Section 2.4 for more details. The first condition in Assumption B(2) is related to the convergence rate of $\widehat{\boldsymbol{\gamma}}_l$, and it ensures that the signals are large enough to be detected. If $C_1 < \kappa_n < C_2$ for some positive constants C_1 and C_2, this condition is simply $\log p < n^{1-c_\beta}/L$ for some small positive constant c_β, which is fulfilled by assumption (3) on p. A few more assumptions on the coefficient functions $\overline{\beta}_j(t)$ will be given in Section 3. The last condition in Assumption B(2) is to ensure that, when the BIC is used as the stopping criterion, our method can deal with ultra-high dimensional cases. For example, if L is taken of the optimal order $n^{1/5}$ then p can be taken as $p = \exp(n^c)$ for any $0 < c \leq 1/5$.

2.4 Forward feature screening

Some of the assumptions we impose in Section 3 may not hold. For example, Assumption B(1) may not hold if some of the irrelevant variables have strong correlation with the true covariates indexed by S_0. Thus, such assumptions may be too restrictive in practice, in particular when p is very large and p_0 is much smaller than p as specified in (3) and (4). In that case, the proposed forward selection procedure may be still used as a forward screening method under certain less restrictive conditions. Then, although some unimportant variables may pass the forward screening, we can utilize some variable selection method to remove them at the next stage. In this section we discuss the details.

Suppose there is a subset of indexes, denoted by $\overline{S}_0$, that contains S_0, and the covariates in $\overline{S}_0$ do not have much correlation with those in $\overline{S}_0^c$. To be clear, we specify the conditions as follows:

(a) $S_0 \subset \overline{S}_0$ and $\#\overline{S}_0 \leq C\#S_0$ for some positive constant C.

(b) $\max_{j \in \overline{S}_0 - S} \|\overline{\beta}_j\|_{L_2} / \max_{j \in \overline{S}_0^c} \|\overline{\beta}_j\|_{L_2} \to \infty$ uniformly for S satisfying $S \subsetneq \overline{S}_0$ and $S_0 \not\subset S$.

(c) Assumption B(2) holds with κ_n replaced with κ_n', where κ_n' is defined by

$$\kappa_n' = \inf_S \max_{j \in S_0 - S} \|\overline{\beta}_j\|_{L_2},$$

with S satisfying the same conditions as in (b).

If we replace conditions B(1) and B(2) with conditions (b) and (c), respectively, and if condition (a) holds, then our procedure given in Section 2.2 can be used as a forward independence screening procedure with an effective stopping rule. That is, it will effectively select all the true covariates indexed by S_0, possibly along with some irrelevant ones from those indexed by $\overline{S}_0 - S_0$. See Proposition 1 given in Section 3 for the theoretical justifications. Those remaining irrelevant covariates will be removed when we apply at the second stage the group SCAD or adaptive group Lasso [5, 12].

3 Assumptions and theoretical properties

In this section, we describe technical assumptions, and we present desirable theoretical properties of the proposed forward procedure in Theorems 1 and 2. Note that we treat the EBIC and the BIC ($\eta = 0$) in a unified way. The proofs are given in Section 5.

First we describe assumptions on the index variable T in the varying coefficient model (1). The following assumption is a standard one when we employ spline estimation.

Assumption T. The index variable T has density function $f_T(t)$ such that $C_{T1} < f_T(t) < C_{T2}$ uniformly in $t \in [0, 1]$, for some positive constants C_{T1} and C_{T2}.

We define some more notation before we state our assumptions on the covariates. Let $\boldsymbol{X}_S$ consist of $\{X_j\}_{j \in S}$ and then $\boldsymbol{X}_S$ is a $\#S$-dimensional random vector. Note that $\boldsymbol{X}_{S(l)}$ is a $(\#S + 1)$-dimensional random vector. For a symmetric matrix $\mathbf{A}$, we denote the maximum and minimum eigenvalues respectively by $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$, and we define $|\mathbf{A}|$ as

$$|\mathbf{A}| = \sup_{|\mathbf{x}|=1} |\mathbf{A}\mathbf{x}| = \max\{|\lambda_{\max}(\mathbf{A})|, |\lambda_{\min}(\mathbf{A})|\}.$$

Assumption X.

X(1) There is a positive constant C_{X1} such that $|X_j| \leq C_{X1}$, $j = 1, \ldots, p$.

X(2) Uniformly in $S \subsetneq S_0$ and $l \in S^c$,

$$C_{X2} \leq \lambda_{\min}(\mathrm{E}\{\boldsymbol{X}_{S(l)}\boldsymbol{X}_{S(l)}^T|T\}) \leq \lambda_{\max}(\mathrm{E}\{\boldsymbol{X}_{S(l)}\boldsymbol{X}_{S(l)}^T|T\}) \leq C_{X3}$$

for some positive constants C_{X2} and C_{X3}.

We use the second assumption X(2) when we evaluate eigenvalues of the matrix $\mathrm{E}\{n^{-1}\boldsymbol{W}_{S(l)}^T\boldsymbol{W}_{S(l)}\}$. We can relax Assumption X(1) slightly by replacing C_{X1} with $C_{X1}(\log n)^{c_X}$ for some positive constant c_X. These are standard assumptions in the variable selection literature.

Assumption **E** below is about the error term ϵ in our varying coefficient model (1). The second condition E(2) requires that ϵ should have the sub-Gaussian property. We use it when we prove the latter half of Theorem 2. This is a standard assumption in the Lasso literature, for example, see [2] and [29].

Assumption E.

E(1) There are positive constants C_{E1} and C_{E2} such that

$$\mathrm{E}\{\exp(C_{E1}|\epsilon|)|X_1, \ldots, X_p, T\} \leq C_{E2}.$$

E(2) There is a positive constant C_{E3} such that $\mathrm{E}\{\exp(u\epsilon)|X_1, \ldots, X_p, T\} \leq \exp(C_{E3}u^2/2)$ for any $u \in \mathbb{R}$.

We need some additional assumptions on the coefficient functions $\overline{\beta}_j$ in the extended marginal model (5) in order to approximate them by the B-spline basis. Note that, in Assumptions B(4)-(5) below, $\overline{\beta}_j \equiv \beta_{0j}$ for all $j \in S_0$ and $\overline{\beta}_j \equiv 0$ for all $j \in S_0^c$ when $S = S_0$.

Assumption B(3)-(5).

B(3) $\kappa_n L^2 \to \infty$ and $\kappa_n = O(1)$, where κ_n is defined in Assumption B(2).

B(4) $\overline{\beta}_j$ is twice continuously differentiable for any $j \in S(l)$ for $S \subset S_0$ and $l \in S^c$.

B(5) There are positive constants C_{B2} and C_{B3} such that $\sum_{j \in S(l)} \|\overline{\beta}_j\|_\infty < C_{B2}$ and $\sum_{j \in S(l)} \|\overline{\beta}_j''\|_\infty < C_{B3}$ uniformly in $S \subset S_0$ and $l \in S^c$.

Theorem 1 given below suggests that the forward selection procedure using criterion (11) can pick up all the relevant covariates in the varying coefficient model (1) when C_{B1} in Assumption B(1) is large enough.

Theorem 1 *Assume that Assumptions* **T**, **X**, *B(1)-(5), and E(1) hold, and define* l^* *as in (11) for any* $S \subsetneq S_0$. *Then, with probability tending to 1, there is a positive constant* C_L *such that*

$$\frac{\left\|\overline{\beta}_{l^*}\right\|_{L_2}}{\max_{j \in S_0 - S} \left\|\overline{\beta}_j\right\|_{L_2}} > C_L$$

uniformly in S, *and thus we have* $l^* \in S_0 - S$ *for any* $S \subsetneq S_0$ *when* C_{B1} *in Assumption B(1) is larger than* $1/C_L$.

Theorem 2 given next implies that the proposed forward procedure will not stop until all of the relevant variables indexed by S_0 have been selected, and it does stop when all the true covariates in model (1) have been selected. Note that in the second result, we have to replace Assumption E(1) with E(2) in order to evaluate a quadratic form of error terms in the proof.

Theorem 2 *Assume that Assumptions* **T**, **X**, *B(1)-(5), and E(1) hold. Then we have the following results.*

(i) For l^* *as in Theorem 1, we have*

$$EBIC(S(l^*)) < EBIC(S)$$

uniformly in $S \subsetneq S_0$, *with probability tending to 1.*

(ii) If we replace Assumption E(1) with Assumption E(2), then we have

$$EBIC(S_0(l)) > EBIC(S_0)$$

uniformly in $l \in S_0^c$, *with probability tending to 1.*

The forward method may also choose some irrelevant covariates if Assumption B(1) fails to hold. In that case, Proposition 1 provides some theoretical results in the setup described in Section 2.4. Note that some conformable changes to Assumptions B(3)-(5) and X(2) and the proofs are needed. See Section 5 for the changes in the proofs.

Proposition 1 *Consider the setup given in Section 2.4. Under the same conditions in Theorem 1 (or Theorem 2), with conformable changes to Assumptions B(3)-(5) and X(2), we have the following results.*

(i) The selected index l^* comes only from $\overline{S}_0$ with probability tending to 1, as in Theorem 1.

(ii) With probability tending to 1, the proposed forward selection procedure continues the feature selection until all the covariates indexed by S_0 are selected, and it stops the selection when all the covariates indexed by S_0 have been selected.

Proposition 1 implies that the proposed forward selection procedure can be used as a forward screening method with an effective stopping rule. Note that, in this setup, we may select some irrelevant covariates from those indexed by $\overline{S}_0 - S_0$. However, the number of potential covariates will be sufficiently reduced after the forward screening stage. Thus, we will be able to remove those remaining irrelevant covariates at the next stage, by using the group SCAD or the adaptive group Lasso [5, 12].

4 Simulation and empirical studies

We carried out two simulation studies and a real data analysis based on the well-known Boston housing data to assess the performance of the proposed forward feature selection method with BIC or EBIC as the stopping criterion. For simplicity, we denote these two variants by fBIC and fEBIC respectively. At the initial step of the forward selection, we let $S_1 = \{0\}$ i.e. we start with the model with only the intercept function. Note that it may happen that the BIC/EBIC drops in one iteration, then increases in the next iteration, and then drops again. To avoid interference caused by such small fluctuations, we continued the fBIC/fEBIC forward selection until the BIC/EBIC continuously increases for five consecutive iterations. The value of the parameter η in the definition (13) of EBIC was taken as $\eta = 1 - \log n/(3 \log p)$, as suggested by [4]. Since the EBIC uses a much larger penalty than the BIC does, it is expected that the fEBIC will select a smaller model than that selected by the fBIC. We could modify the penalty term by adjusting the degrees of freedom or change the value of η to a smaller one, but it becomes complicated.

In the simulation studies, we generated data from the two varying coefficient models studied by [10]. Following the paper, we used the cubic B-spline with $L = 7$, we set the sample size and the number of covariates as $n = 400$ and $p = 1000$ respectively, and we repeated each of the simulation configuration for $N = 200$ times.

Table 1: Correlations between the covariates X_j's and the index variable T.

$[t_1, t_2]$	$[0,0]$	$[2,0]$	$[3,0]$	$[2,1]$	$[3,1]$	$[3,2]$
$\text{corr}(X_j, X_k)$	0	0.25	0.43	0.25	0.43	0.43
$\text{corr}(X_j, W)$	0	0	0	0.36	0.46	0.59

4.1 Comparison of fBIC and fEBIC

In this section, we compare the finite sample performance of the fBIC and the fEBIC using the two varying coefficient models studied by [10].

Example 1 *Following Example 3 of [10], we generated N samples from the following varying coefficient model:*

$$Y = 2 \cdot X_1 + 3\,T \cdot X_2 + (T+1)^2 \cdot X_3 + \frac{4\sin(2\pi T)}{2 - \sin(2\pi T)} \cdot X_4 + \epsilon,$$

where $X_j = (Z_j + t_1 U_1)/(1 + t_1), j = 1, 2, \cdots, p$, and $T = (U_2 + t_2 U_1)/(1 + t_2)$, with $Z_1, Z_2, \cdots, Z_p \overset{i.i.d}{\sim} N(0,1), U_1, U_2 \overset{i.i.d}{\sim} U(0,1)$, and $\epsilon \sim N(0,1)$ being all mutually independent with each other.

In this example, the number of true covariates p_0 is four. The tuning parameters t_1 and t_2 are used to control the correlations between the covariates $X_j, j = 1, 2, \cdots, p$ and the index covariate T. It is easy to show that $\text{corr}(X_j, X_k) = t_1^2/(12 + t_1^2)$ for any $j \neq k$, and $\text{corr}(X_j, T) = t_1 t_2/[(12 + t_1^2)(1 + t_2^2)]^{1/2}$ independent of j. Table 1 lists the values of the tuning parameters $[t_1, t_2]$ which define six cases of the correlations between the covariates X_j's and the index covariate T. The first case is associated with the situation when the X_j's are uncorrelated while they are uncorrelated with T. The second and third cases are associated with those situations when the X_j's are increasingly correlated but they are uncorrelated with T. The last three cases are associated with those situations when the X_j's are increasingly correlated and the correlations between the X_j's and T are also increasing. These six cases allow us to compare the performance of the fBIC and fEBIC procedures effectively, In the next section, we will also use them to compare the performance of the fBIC with those procedures proposed and studied by [10].

Figure 1 depicts the boxplots of the model sizes selected by the fBIC and the fEBIC in the six correlation cases. It is seen that in all the six cases, the fBIC performs very well in terms of correctly selecting the right model except that it occasionally selects a

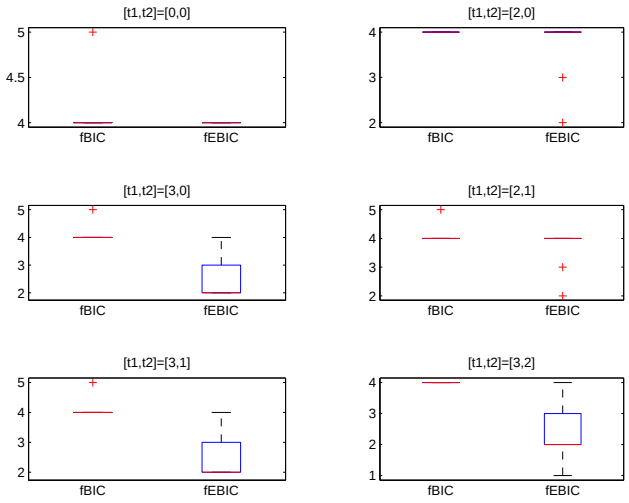

Figure 1: *Boxplots of the model sizes selected by the fBIC and fEBIC for the varying coefficient model in Example 1*

model with one extra covariate out of the 200 runs. However, generally speaking the fEBIC selects a smaller model as compared to the true model, and it selects all of the four true covariates most of the time only when the correlations between the X_j's and T are relatively small. As the correlations between the X_j's or the correlations between the X_j's and T increase, the performance of fEBIC becomes worse and it selects a much smaller model than the correct one most of the time.

The varying coefficient model in Example 1 has only four true underlying covariates. In the varying coefficient model defined in the following example, there are eight true underlying covariates.

Example 2 *Following Example 4 of [10], we generated N samples from the following varying coefficient model:*

$$\begin{aligned} Y = \quad & 3\,T \cdot X_1 + (T+1)^2 \cdot X_2 + (T-2)^3 \cdot X_3 + 3(\sin(2\pi T)) \cdot X_4 \\ & + \exp(T) \cdot X_5 + 2 \cdot X_6 + 2 \cdot X_7 + 3\sqrt{T} \cdot X_8 + \epsilon, \end{aligned}$$

while T, $\mathbf{X}$, Y and ϵ were generated in the same way as described in Example 1.

Figure 2 shows the boxplots of the model sizes selected by fBIC and fEBIC in the six correlation cases given in Table 1, when the data came from the varying coefficient model defined in Example 2. Again, we observe that in all these six cases, the fBIC performs very well in terms of correctly selecting the right model except that it occasionally

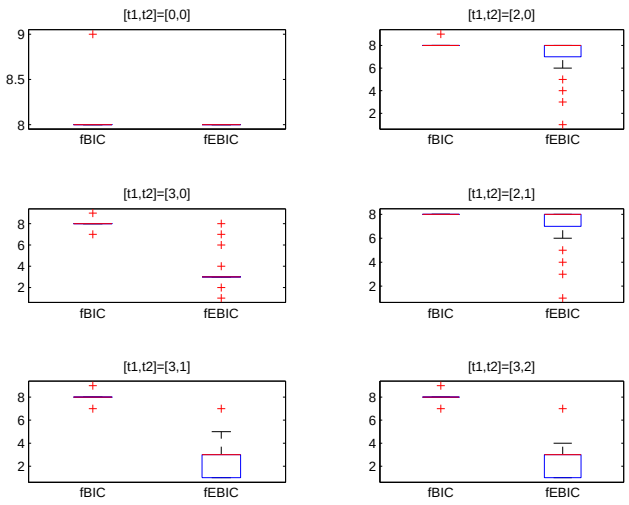

Figure 2: *Boxplots of the model sizes selected by the fBIC and fEBIC for the varying coefficient model in Example 2*

selects a model with one extra or one less covariate out of the 200 runs. However, the fEBIC selects a smaller model in general, and it selects the right model only when the correlations between X_j's and T are relatively small. Similar to Example 1, when the correlations between X_j's or in the correlations between X_j's and T increase, the performance of fEBIC becomes worse and it selects a much smaller model than the right model most of the time.

From the above two examples, we see that the fBIC consistently outperforms the fEBIC substantially. It appears that when a forward selection procedure is used in the considered context, the BIC-based stopping rule is better than the one using EBIC, since the EBIC penalizes the introduction of a new covariate too much and as a result it stops too early. This may seem to contradict with the rational behind the original EBIC designed for linear models. But, for varying coefficient models the degrees of freedom in the definition of EBIC increases much faster when more variables are introduced to the model. Note also the original EBIC is introduced for model selection, not forward selection. Following the observation that fBIC performs very well numerically and the fact that η disappears from it, we prefer the fBIC to the fEBIC for the studied problem.

4.2 Comparison with the approaches of Fan, Ma, and Dai (2014)

In this section, we compare the performance of the fBIC with that of the conditional-INIS and the greedy-INIS approaches introduced by [10]. We consider exactly the same

Table 2: Average numbers of true positive (TP) and false positive (FP), and prediction error (PE) over 200 repetitions and their robust standard deviations (in parentheses) for the conditional-INIS, greedy-INIS and fBIC approaches under the varying coefficient model defined in Example 1.

$[t_1, t_2]$	SNR	Conditional-INIS			Greedy-INIS			fBIC		
		TP	FP	PE	TP	FP	PE	TP	FP	PE
$[0, 0]$	16.85	4	0.54	1.10	4	13.01	1.41	4	0	0.95
		(0)	(0.75)	(0.05)	(0)	(3.73)	(0.17)	(0)	(0)	(0.04)
$[2, 0]$	3.66	4	0.20	0.78	4	0.41	1.10	4	0.01	1.12
		(0)	(0)	(0.06)	(0)	(0)	(0.05)	(0)	(0)	(0.05)
$[3, 0]$	3.32	4	0.19	1.03	3.99	0.57	1.22	4	0.01	1.20
		(0)	(0)	(0.06)	(0)	(0)	(0.07)	(0)	(0)	(0.04)
$[2, 1]$	3.21	3.97	0.26	1.27	3.90	1.14	1.63	4	0	1.20
		(0)	(0)	(0.24)	(0)	(0)	(0.41)	(0)	(0)	(0.07)
$[3, 1]$	2.81	3.95	0.31	1.30	3.77	0.27	1.29	3.99	0	1.18
		(0)	(0.75)	(0.12)	(0)	(0)	(0.17)	(0)	(0)	(0.07)

simulation setups as their Examples 3 and 4 and adopt their simulation results. Following [10], we report the average numbers of true positive (TP) and false positive (FP) selections, the prediction error (PE), and their robust standard deviations for all the three procedures under consideration, where the prediction error is the mean squared error calculated on a test dataset of size $n/2 = 200$ randomly generated from the same model. The signal-to-noise-ratio, denoted by SNR and defined as $\mathrm{Var}(\beta^T(T)\mathbf{X})/\mathrm{Var}(\epsilon)$, is also reported as it is an important measure of the complexity of the varying coefficient model associated with the tuning parameters $[t_1, t_2]$.

Table 2 displays the simulation results under the varying coefficient model defined in Example 1. We can see that the fBIC in general outperforms both the conditional-INIS and the greedy-INIS approaches in terms of the values of TP, FP, and PE. In the first three cases where X_j's and T are uncorrelated, all the three procedures are comparable in terms of selecting correctly all of the true covariates, but the fBIC selects fewer false covariates than the other two competitors and the fBIC also has smaller values of PE in general. In the latter two cases where X_j's and T are correlated, the performance of the conditional-INIS and greedy-INIS approaches become worse while the performance of fBIC is still good in terms of the values of TP, FP, and PE. The good performance of fBIC is consistent with what we observed from Figure 1.

Table 3: The same as that of Table 2 but now under the varying coefficient model defined in Example 2.

$[t_1, t_2]$	SNR	Conditional-INIS			Greedy-INIS			fBIC		
		TP	FP	PE	TP	FP	PE	TP	FP	PE
$[0, 0]$	47.68	8	0.21	1.24	8	10.71	1.57	8	0.02	1.22
		(0)	(0)	(0.09)	(0)	(3.73)	(0.20)	(0)	(0)	(0.09)
$[2, 0]$	9.40	8	0.13	1.17	8	0.60	1.16	8	0	1.20
		(0)	(0)	(0.09)	(0)	(0)	(0.10)	(0)	(0)	(0.08)
$[3, 0]$	8.18	7.90	0.10	1.21	7.98	0.71	1.29	7.99	0.03	1.18
		(0)	(0)	(0.12)	(0)	(0)	(0.10)	(0)	(0)	(0.11)
$[2, 1]$	8.62	7.80	0.20	2.16	7.55	0.26	2.26	8	0.01	2.55
		(0)	(0)	(0.58)	(0.75)	(0)	(0.70)	(0)	(0)	(0.64)
$[3, 1]$	7.61	7.75	0.18	1.65	7.35	0.28	1.84	7.96	0.02	1.37
		(0)	(0)	(0.26)	(0.75)	(0)	(0.42)	(0)	(0)	(0.22)

Table 3 displays the simulation results under the varying coefficient model defined in Example 2. Similarly, it is seen that fBIC in general outperforms the conditional-INIS and greedy-INIS approaches. Along with increases in the correlations between X_j's and the correlations between X_j's and T, the performance of the conditional-INIS and greedy-INIS approaches become worse very quickly while the performance of fBIC becomes worse much more slowly. The good performance of the fBIC is consistent with what we observed from Figure 2.

4.3 Applications to the Boston housing data

Following [10], we applied the fBIC approach to the well-known Boston housing dataset (Harrison and Rubinfeld 1978) whose description can be found in the manual of R package *mlbench*. The dataset contains 506 census tracts of Boston from the 1970 census with 13 covariates. The housing value equation obtained in the literature, as reported by [14], can be written as

$$\begin{aligned} \log(MV) &= \beta_0 + \beta_1 RM^2 + \beta_2 AGE + \beta_3 \log(DIS) \\ &\quad + \beta_4 \log(RAD) + \beta_5 TAX + \beta_6 PTRATIO \\ &\quad + \beta_7 (B - 0.63)^2 + \beta_8 \log(LSTAT) + \beta_9 CRIM \\ &\quad + \beta_{10} ZN + \beta_{11} INDUS + \beta_{12} CHAS \\ &\quad + \beta_{13} NOX^2 + \epsilon, \end{aligned} \tag{14}$$

where the dependent variable MV is the median value of owner-occupied homes, and the independent covariates are quantified measurements of its neighborhood. To adopt a varying coefficient model for the Boston housing data, [10] took the covariate $\log(DIS)$, the weighted distance to five employment centers in the Boston region, as the index variable T and replaced the constant coefficients β_j in (14) with the varying coefficients $\beta_j(T)$. This allows us to examine how the weighted distance to the five employment centers interacts with the other covariates. It seems reasonable to assume that the impacts of the other covariates on housing price change with this distance. Using the conditional-INIS approach, [10] obtained the following varying coefficient submodel:

$$\begin{aligned} \log(MV) &= \beta_0(T) + \beta_1(T)RM^2 + \beta_2(T)AGE + \beta_5(T)TAX \\ &\quad + \beta_7(T)(B - 0.63)^2 + \beta_9(T)CRIM + \epsilon. \end{aligned} \tag{15}$$

By the fBIC approach, we obtained the following varying coefficient submodel:

$$\begin{aligned} \log(MV) &= \beta_0(T) + \beta_1(T)RM^2 + \beta_2(T)AGE \\ &\quad + \beta_6(T)PTRATIO + \beta_7(T)(B - 0.63)^2 \\ &\quad + \beta_8(T)\log(LSTAT) + \beta_9(T)CRIM \\ &\quad + \beta_{13}(T)NOX^2 + \epsilon. \end{aligned} \tag{16}$$

It is interesting to compare the two varying coefficient submodels (15) and (16) selected by the conditional-INIS approach of [10] and the fBIC procedure respectively. We can see that model (16) does not introduce the covariate TAX which is introduced in model (15), while it includes three other covariates $PTRATIO, \log(LSTAT)$, and NOX^2 which are not present in model (15). Notice that the covariate $PTRATIO$ denotes the pupil-teacher ratio by the town school district, and a lower ratio indicates each student receives more individual attention. It is reasonable that parents usually want to buy houses near good schools which tend to have smaller values of $PTRATIO$. Therefore, it is expected that $PTRATIO$ should have important negative impact on housing values. Notice also that the covariate $LSTAT$ is the proportion of the population that is of lower status. It is natural that a larger proportion of poor people in a region often means lower average housing prices in that region. Therefore, $LSTAT$ should have important negative impact on the housing values. Finally notice that the covariate NOX is a measure for air pollution level, and it generally has a negative impact on the housing values since people usually want to live in a region where there is less air pollution. In summary, introduction of these three covariates in the model (16) sounds reasonable. In fact the correlations between the covariates $PTRATIO, \log(LSTAT)$,

and NOX^2 and the response $\log(MV)$ are $-0.5017, -0.8230$, and -0.4965 respectively. As for the covariate TAX, there is no doubt that it is an important covariate which may have important negative impact on the housing evaluation; in fact, the correlation between TAX and $\log(MV)$ is -0.5615. On the other hand, it also has strong correlations with $PTRATIO, \log(LSTAT)$, and NOX^2, which are $0.5224, 0.4609$, and 0.6415 respectively. Therefore, with introduction of $PTRATIO, \log(LSTAT)$, and NOX^2 in the model already, the effect of TAX on $\log(MV)$ may have been represented by that of $PTRATIO, \log(LSTAT)$, and NOX^2.

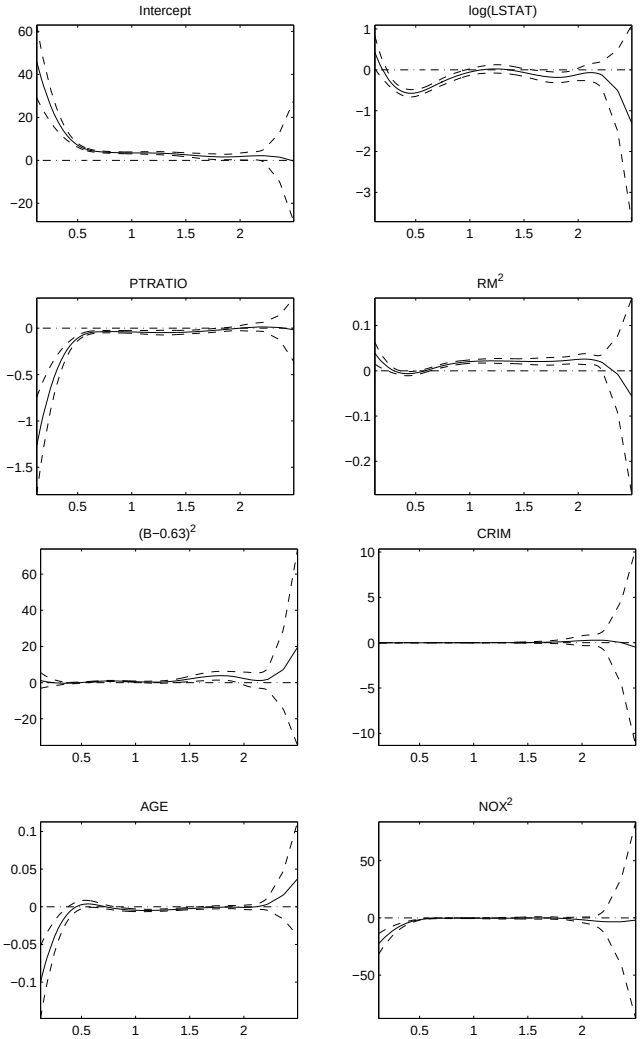

Figure 3: *Fitted coefficient functions (solid) with approximate 95% confidence bands (dashed) for the Boston housing data. Cubic B-splines with the number of basis functions, $L_n = 7$, selected by fBIC, were used.*

Figure 3 plots the fitted coefficient functions $\beta_j(T)$'s, along with the corresponding

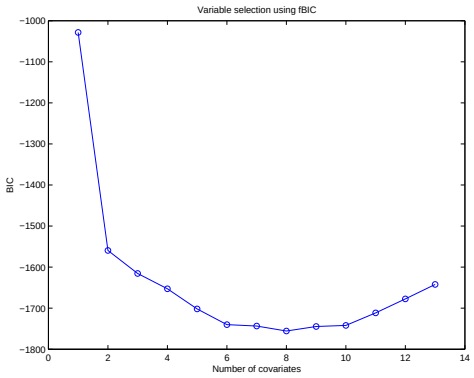

Figure 4: *Variable selection using fBIC for the Boston housing data.*

approximate 95% confidence bands, according to the order in which they were selected by the fBIC, that is, the covariate $\log(LSTAT)$ was first selected, followed by the covariate $PTRATIO$, and then RM^2, etc. Figure 4 displays the BIC curve for the forward variable selection when applied to the Boston housing data. From Figure 3, it is seen that the introduction of $\log(LSTAT)$ in the model (16) at the first selection step indicates that it has the most important impact on the housing values in the Boston regions under consideration, and the socioeconomic status distinctions mean more in the upper brackets of the society than in the lower classes. The associated coefficient curve shows that the impact of $\log(LSTAT)$ on housing values is generally negative as expected, especially when the regions are near the five employment centers. The effect at both ends are not significant and may be due to boundary effect of B-spline smoothing when less data are available. The introduction of $PTRATIO$ at the second step indicates that this covariate also has important impact on the housing value. The associated coefficient curve shows that the impact is negative, especially at those regions near the five employment centers. The covariate RM is the third covariate introduced in the model (16), and it is the average number of rooms in owner units, which represents the size of a house. As expected, this covariate has positive impact on the housing value. The impacts of the other four selected covariates on housing values can be analyzed and interpreted similarly; see [10] and [14] for more details.

The Boston housing data set has only twelve covariates under consideration with $\log(DIS)$ as the index covariate. It can not be regarded as a real high-dimensional data example. To overcome this difficulty, [10] extended the Boston housing data via introducing the following artificial covariates:

$$X_j = \frac{Z_j + 2U}{3}, j = 13, 14, \cdots, 1000,$$

Table 4: Prediction error (PE), model size (MS), and selected noise variables (SNV) over 100 repetitions and their robust standard deviations (in parentheses) for the conditional-INIS, greedy-INIS, fBIC, modified fBIC approaches.

Approach	PE	MS	SNV
Conditional-INIS	0.046 (0.048)	5.55 (0.75)	0 (0)
Greedy-INIS	0.048 (0.020)	4.80 (1.49)	0.01 (0)
fBIC	0.083 (0.033)	8.60 (2.24)	2.16 (1.49)
Modified fBIC	0.049 (0.019)	7.28 (1.49)	0.63 (0.75)
fBIC-SCAD	0.062 (0.023)	7.00 (1.49)	1.89 (1.49)

where $Z_j, j = 13, \cdots, 1000 \overset{i.i.d}{\sim} N(0,1)$ and $U \sim U[0,1]$ are independent. They randomly selected $n = 406$ observations as the training set and applied their conditional-INIS and greedy-INIS approaches to select the models, and then computed the associated prediction mean squared error (PE) on the rest 100 observations. This process was repeated $N = 100$ times and they reported the average prediction error and model size, and their robust standard deviations as in Table 4. We repeated the above process with the fBIC approach and the results are also displayed in the table. It turns out that the fBIC approach selects a few artificial covariates. This is consistent with those observed in Figures 1 and 2.

To overcome this difficulty, we can first rank the covariates according to the BIC values of their corresponding marginal models, and then apply the fBIC approach to the data with the first fifty covariates, say. The associated approach is called the modified fBIC approach. Since the dimensionality becomes smaller and it is expected that the fBIC approach will perform better in this case. The results presented in Table 4 indicate that the average model size selected by the modified fBIC approach is indeed better than that selected by the fBIC approach, and it is about the same as that of model (16) which is selected when there are only twelve covariates involved. In addition, the PE and SNV values show that the modified fBIC approach improves on the fBIC approach substantially and that it is comparable with the Conditional-INIS and the Greedy-INIS. Alternatively, as mentioned in Section 2.4, we may apply the fBIC approach first and then apply the group SCAD to further remove those unwanted covariates. The resulting approach may be termed as the fBIC-SCAD approach, and the associated simulation results are listed at the last row of Table 4. The results show that applying group SCAD indeed improves the performance of the fBIC approach.

From this example, it is seen that the fBIC approach or its modified version is very useful in scientific discoveries based on high-dimensional data with complex structure. It can select a parsimonious close-to-truth model, and can reveal interesting relationship between the response variable and the important covariates.

5 Proofs

First, we define some notation related to the approximate regression models (6) and (7). Let

$$\begin{aligned} D_{lSn} &= n^{-1}\boldsymbol{W}_{S(l)}^T\boldsymbol{W}_{S(l)} \quad \text{and} \quad D_{lS} = \mathrm{E}\{D_{lS_n}\}, \\ d_{lS_n} &= n^{-1}\boldsymbol{W}_{S(l)}^T\boldsymbol{Y} \quad \text{and} \quad d_{lS} = \mathrm{E}\{d_{lS_n}\}, \quad \text{and} \\ \Delta_{lSn} &= D_{lSn}^{-1}d_{lSn} - D_{lS}^{-1}d_{lS}\,. \end{aligned}$$

Then, the parameter vector $\overline{\boldsymbol{\gamma}}_l$ in model (7) can be expressed as $\overline{\boldsymbol{\gamma}}_l = (\mathbf{0}_L, \ldots, \mathbf{0}_L, \mathbf{I}_L)D_{lS}^{-1}d_{lS}$, where $\mathbf{0}_L$ denotes the $L \times L$ zero matrix and $\mathbf{I}_L$ is the L-dimensional identity matrix.

Before we prove Theorems 1 and 2, we present Lemmas 1-3. We verify these lemmas at the end of this section. In Lemma 1 we evaluate the minimum and maximum eigenvalues of some matrices.

Lemma 1 *Assume that Assumptions* **T**, **X**, *and E(1) hold. Then, with probability tending to 1, there are positive constants* M_{11}, M_{12}, M_{13}, *and* M_{14} *such that*

$$L^{-1}M_{11} \le \lambda_{\min}(D_{lSn}) \le \lambda_{\max}(D_{lSn}) \le L^{-1}M_{12}$$

and

$$L^{-1}M_{13} \le \lambda_{\min}(n^{-1}\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}) \le \lambda_{\max}(n^{-1}\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}) \le L^{-1}M_{14}$$

uniformly in $S \subsetneq S_0$ *and* $l \in S^c$.

Lemma 2 is about the relationship between β_l and $\overline{\boldsymbol{\gamma}}_l$ in the extended marginal models (5) and (6).

Lemma 2 *Assume that Assumptions* **T**, **X**, *and B(4)-(5) hold. Then there are positive constants* M_{21} *and* M_{22} *such that*

$$M_{21}\sqrt{L}\big(\|\overline{\beta}_l\|_{L_2} - O(L^{-2})\big) \le |\overline{\boldsymbol{\gamma}}_l| \le M_{22}\sqrt{L}\big(\|\overline{\beta}_l\|_{L_2} + O(L^{-2})\big)$$

uniformly in $S \subsetneq S_0$ *and* $l \in S^c$.

We use Lemma 3 to evaluate the estimation error for $\overline{\boldsymbol{\gamma}}_j$, $j \in S(l)$, in model (6).

Lemma 3 *Assume that Assumptions* **T**, **X**, *and B(4)-(5) hold. Then, for any* $\delta > 0$, *there are positive constants* M_{31}, M_{32}, M_{33}, *and* M_{34} *such that*

$$|\Delta_{lSn}| \leq M_{31} L^{3/2} p_0^{3/2} \delta / n$$

uniformly in $S \subsetneq S_0$ *and* $l \in S^c$, *with probability*

$$1 - M_{32}\, p_0^2\, L \exp\Big\{ - \frac{\delta^2}{M_{33} n L^{-1} + M_{34}\delta} + \log p + p_0 \log 2 \Big\}.$$

5.1 Proofs of Theorems 1 and 2, and Proposition 1

Now we prove Theorems 1 and 2 by employing Lemmas 1-3.

Proof of Theorem 1. Consider the case that $S \subsetneq S_0$ and $l \in S^c$. Note we can write

$$\widehat{\boldsymbol{\gamma}}_l = \overline{\boldsymbol{\gamma}}_l + (\mathbf{0}_L, \ldots, \mathbf{0}_L, \mathbf{I}_L)\Delta_{lSn}. \tag{17}$$

Lemma 1 implies we should deal with Δ_{lSn} on the right-hand side of (17) when we evaluate $\widehat{\sigma}_S^2 - \widehat{\sigma}_{S(l)}^2$ given in equation (10). For this purpose, Assumption B(2) suggests that we should take δ in Lemma 3 as $\delta = n^{1-c_\beta/4}\kappa_n / L$ tending to ∞. Recall the definition of κ_n in Assumption B(2). Then we have that

$$\frac{\sqrt{L}\kappa_n}{L^{3/2} p_0^{3/2} \delta / n} = \frac{n^{c_\beta/4}}{p_0^{3/2}} \to \infty \tag{18}$$

and

$$\begin{aligned} & p_0^2 L \exp\Big\{ - \frac{1}{2M_{33}} \frac{\delta^2}{nL^{-1}} + \log p + p_0 \log 2 \Big\} \\ & = p_0^2 L \exp\Big\{ - (2M_{33})^{-1} n^{1-c_\beta/2} \kappa_n^2 L^{-1} + \log p + p_0 \log 2 \Big\} \\ & < p_0^2 L \exp\Big\{ - (2M_{33})^{-1} n^{c_\beta/2} \log p + \log p + p_0 \log 2 \Big\} \to 0. \end{aligned} \tag{19}$$

By (18), (19), and Lemma 3, $(\mathbf{0}_L, \ldots, \mathbf{0}_L, \mathbf{I}_L)\Delta_{lSn}$ is negligible compared to $\boldsymbol{\gamma}_l$ on the right-hand side of (17), with probability tending to 1. Therefore Lemmas 1 and 2 and Assumption B(3) imply that we should focus on $\sqrt{L}\|\beta_l\|$ in evaluating $\widehat{\sigma}_{S(l)}^2$ in (10). Hence the desired result follows from Assumption B(1). □

Proof of Theorem 2. To prove result (i), we evaluate

$$\mathrm{EBIC}(S) - \mathrm{EBIC}(S(l)) = n \log \Big(\frac{n\widehat{\sigma}_S^2}{n\widehat{\sigma}_{S(l)}^2} \Big) - L(\log n + 2\eta \log p).$$

Since

$$n\widehat{\sigma}_S^2 - n\widehat{\sigma}_{S(l)}^2 = (\widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{Y}}_S)^T (\widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{W}}_{lS})^{-1} (\widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{Y}}_S) = \widehat{\boldsymbol{\gamma}}_l^T \widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{W}}_{lS} \widehat{\boldsymbol{\gamma}}_l,$$

we have

$$\frac{n\widehat{\sigma}_S^2}{n\widehat{\sigma}_{S(l)}^2} \geq 1 + (n^{-1}\boldsymbol{Y}^T\boldsymbol{Y})^{-1} \widehat{\boldsymbol{\gamma}}_l^T \Big(\frac{1}{n} \widetilde{\boldsymbol{W}}_{lS}^T \widetilde{\boldsymbol{W}}_{lS} \Big) \widehat{\boldsymbol{\gamma}}_l. \tag{20}$$

Then Lemma 1 and (20) imply that we have for some positive C,

$$\mathrm{EBIC}(S) - \mathrm{EBIC}(S(l)) \geq CnL^{-1} |\widehat{\boldsymbol{\gamma}}_l|^2 - L(\log n + 2\eta \log p) \tag{21}$$

uniformly in $S \subsetneq S_0$ and $l \in S^c$, with probability tending to 1. Here we use the fact that $L^{-1}|\widehat{\boldsymbol{\gamma}}_l|^2$ is uniformly bounded with probability tending to 1. Then as in the proof of Theorem 1, we should consider $\sqrt{L}\|\overline{\beta}_j\|$ in evaluating the right-hand side of (21). Since Assumption B(2) implies that

$$\frac{nL^{-1}(\sqrt{L}\kappa_n)^2}{L(\log n + 2\eta \log p)} = \frac{n\kappa_n^2}{L(\log n + 2\eta \log p)} \to \infty,$$

we have from (21) that

$$\mathrm{EBIC}(S) - \mathrm{EBIC}(S(l)) > 0$$

uniformly in $S \subsetneq S_0$ and $l \in S^c$ satisfying $\|\overline{\beta}_l\|_{L_2} \big/ \max_{j \in S_0 - S} \|\beta_j\|_{L_2} > C_L$, with probability tending to 1. Hence the proof of result (i) is complete.

To prove result (ii), recall that we replace Assumption E(1) with Assumption E(2). We should evaluate

$$\begin{aligned} &\mathrm{EBIC}(S_0(l)) - \mathrm{EBIC}(S_0) \\ &= \; n \log \Big\{ 1 - \frac{\boldsymbol{Y}^T \widetilde{\boldsymbol{W}}_{lS_0} (\widetilde{\boldsymbol{W}}_{lS_0}^T \widetilde{\boldsymbol{W}}_{lS_0})^{-1} \widetilde{\boldsymbol{W}}_{lS_0}^T \boldsymbol{Y}}{n\widehat{\sigma}_{S_0}^2} \Big\} + L(\log n + 2\eta \log p) \end{aligned} \tag{22}$$

for $l \in S_0^c$. It is easy to prove that $\widehat{\sigma}_{S_0}^2$ converges to $\mathrm{E}\{\epsilon^2\}$ in probability and the details are omitted. We denote $\widetilde{\boldsymbol{W}}_{lS_0} (\widetilde{\boldsymbol{W}}_{lS_0}^T \widetilde{\boldsymbol{W}}_{lS_0})^{-1} \widetilde{\boldsymbol{W}}_{lS_0}^T$ by $\widetilde{\mathbf{P}}_{lS_0}$, which is an orthogonal projection matrix. Thus, from (22) we have for some positive C,

$$\mathrm{EBIC}(S_0(l)) - \mathrm{EBIC}(S_0) \geq -\frac{C}{\mathrm{E}\{\epsilon^2\}} \boldsymbol{Y}^T \widetilde{\mathbf{P}}_{lS_0} \boldsymbol{Y} + L(\log n + 2\eta \log p) \tag{23}$$

uniformly in $l \in S_0^c$, with probability tending to 1.

Now we evaluate $\boldsymbol{Y}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{Y}$ on the right-hand side of (23). From the definition of $\widetilde{\boldsymbol{W}}_{lS_0}$, we have

$$\boldsymbol{Y}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{Y} = (\boldsymbol{Y} - \boldsymbol{W}_{S_0}\boldsymbol{\gamma}_{S_0})^T\widetilde{\mathbf{P}}_{lS_0}(\boldsymbol{Y} - \boldsymbol{W}_{S_0}\boldsymbol{\gamma}_{S_0})$$

for any $\boldsymbol{\gamma}_{S_0} \in \mathbb{R}^{L\#S_0}$. Therefore we obtain

$$\boldsymbol{Y}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{Y} \leq \boldsymbol{\epsilon}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{\epsilon} + |\boldsymbol{b}|^2$$

where $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_n)^T$ and $\boldsymbol{b}$ is some n-dimensional vector of spline approximation errors satisfying $|\boldsymbol{b}|^2 = O(nL^{-4})$ uniformly in $l \in S_0^c$. By applying Proposition 3 of [33], we obtain

$$\mathrm{P}\Big(\frac{\boldsymbol{\epsilon}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{\epsilon}}{LC_{E2}} \geq \frac{1+x}{\{1-2/(e^{x/2}\sqrt{1+x}-1)\}_+^2}\Big) \leq \exp(-Lx/2)(1+x)^{L/2}, \tag{24}$$

where $\{x\}_+ = \max\{0, x\}$. We take $x = \log(p^{2\eta}n)a_n/2$ with a_n tending to 0 sufficiently slowly. Then from the above inequality, we have $\boldsymbol{\epsilon}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{\epsilon} = o_p(L\log(p^{2\eta}n))$ uniformly in $l \in S_0^c$. Thus we have

$$\boldsymbol{Y}^T\widetilde{\mathbf{P}}_{lS_0}\boldsymbol{Y} = O(nL^{-4}) + o_p(L\log(p^{2\eta}n)) \tag{25}$$

uniformly in $l \in S_0^c$. Hence the desired result follows from (23), (25), and the assumption that $L = c_L n^{\kappa_L}$ with $\kappa_L \geq 1/5$. Note that, here we use the condition that $L\log n/\log p \to \infty$ when $\eta = 0$, which is stated in Assumption B(2). □

Proof of Proposition 1. The first result follows from almost the same arguments as in the proof of Theorem 1, thus we omit the proof. We just comment on proof of the second one, which corresponds to result (ii) of Theorem 2. We should deal with S such that $S_0 \subset S \subset \overline{S}_0$ in the proof. Then we replace $\widehat{\sigma}^2_{S_0}$ in (22) with $\widehat{\sigma}^2_S$ and replace $\widetilde{\mathbf{P}}_{lS_0}$ with $\widetilde{\mathbf{P}}_{lS}$ everywhere. Nevertheless, we still have $\boldsymbol{\epsilon}^T\widetilde{\mathbf{P}}_{lS}\boldsymbol{\epsilon} = o_p(L\log(p^{2\eta}n))$ uniformly in S and $l \in S^c$ by exploiting (24). There is no change about the B-spline approximation. Thus we obtain the version of (23) and (25) with S_0 replaced by S, and the modified (23) and (25) hold uniformly in S. Hence the latter half of Proposition 1 is established. Note that some minor conformable changes to the assumptions are necessary.

5.2 Proofs of lemmas

We use the following inequalities in the proofs of Lemmas 1-2.

$$\frac{C_{S1}}{L} \leq \lambda_{\min}(\mathrm{E}\{\mathbf{B}(T)\mathbf{B}(T)^T\}) \leq \lambda_{\max}(\mathrm{E}\{\mathbf{B}(T)\mathbf{B}(T)^T\}) \leq \frac{C_{S2}}{L}, \tag{26}$$

where C_{S1} and C_{S2} are positive constants independent of L. See [15] for the proof of (26).

Proof of Lemma 1. Write

$$n^{-1}D_{lSn} = n^{-1}\sum_{i=1}^{n}(\boldsymbol{X}_{iS(l)}\boldsymbol{X}_{iS(l)}^T) \otimes (\mathbf{B}(T_i)\mathbf{B}(T_i)^T), \tag{27}$$

where $\boldsymbol{X}_{iS(l)}$ is the ith sample version of $\boldsymbol{X}_{S(l)}$ and $\otimes$ is the kronecker product. Note that (26), (27), and Assumption X(2) imply that, for any $\delta > 0$,

$$\frac{C_1}{L} \leq \lambda_{\min}(D_{lS}) \leq \lambda_{\max}(D_{lS}) \leq \frac{C_2}{L}. \tag{28}$$

for some positive C_1 and C_2. In addition, by exploiting the band-diagonal property of D_{lSn} and D_{lS} and an exponential inequality, we can demonstrate that

$$|D_{lSn} - D_{lS}| \leq n^{-1}\delta p_0 \tag{29}$$

uniformly in $S \subsetneq S_0$ and $l \in S^c$ with probability

$$1 - C_3 p_0^2 L \exp\{-\delta^2(C_4 n L^{-1} + C_5\delta)^{-1}\} \times p\exp(p_0 \log 2), \tag{30}$$

where C_3, C_4, and C_5 are positive constants independent of p_0, L, n, p, and δ. When we take $\delta = n^{1-c_\beta/4}L^{-1}$, the probability in (30) tends to 0 and the former result follows since $\delta p_0/n = p_0 n^{-c_\beta/4}/L = o(L^{-1})$. The latter result follows from the following relationship between D_{lSn}^{-1} and $n^{-1}\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}$:

$$D_{lSn}^{-1} = \begin{pmatrix} * & * \\ * & \left(n^{-1}\widetilde{\boldsymbol{W}}_{lS}^T\widetilde{\boldsymbol{W}}_{lS}\right)^{-1} \end{pmatrix}.$$

□

Proof of Lemma 2. Let $\{b_j\}_{j\in S(l)}$ be a set of square integrable functions on $[0,1]$. Then Assumption X(2) implies that

$$C_{X2}\sum_{j\in S(l)}\|b_j(T)\|^2 \leq \|\sum_{j\in S(l)}X_j b_j(T)\|^2 \leq C_{X3}\sum_{j\in S(l)}\|b_j(T)\|^2. \tag{31}$$

Besides, Assumption **T** implies

$$C_{T1}\|b\|_{L_2}^2 \leq \|b(T)\|^2 \leq C_{T2}\|b\|_{L_2}^2 \tag{32}$$

for any square integrable function b. In addition, due to Assumptions B(4) and B(5), we can choose some positive constant C_1 and a set of L-dimensional vectors $\{\tilde{\boldsymbol{\gamma}}_j\}_{j\in S(l)}$ such that

$$\sum_{j\in S(l)}\|\overline{\beta}_j - \tilde{\boldsymbol{\gamma}}_j^T\mathbf{B}\|_\infty \leq C_1 L^{-2}, \tag{33}$$

where C_1 depends only on the assumptions.

By exploiting (31)-(33), we obtain

$$\begin{aligned} & C_{X2} \sum_{j \in S(l)} \|\overline{\beta}_j(T) - \overline{\boldsymbol{\gamma}}_j^T \mathbf{B}(T)\|^2 \\ & \leq \ \| \sum_{j \in S(l)} (\overline{\beta}_j(T) - \overline{\boldsymbol{\gamma}}_j^T \mathbf{B}(T)) X_j \|^2 \ \leq \ \| \sum_{j \in S(l)} (\overline{\beta}_j(T) - \tilde{\boldsymbol{\gamma}}_j^T \mathbf{B}(T)) X_j \|^2 \\ & \leq \ \sum_{j \in S(l)} \|\overline{\beta}_j(T) - \tilde{\boldsymbol{\gamma}}_j^T \mathbf{B}(T)\|^2 \ \leq \ C_{X3} \sum_{j \in S(l)} \|\overline{\beta}_j - \tilde{\boldsymbol{\gamma}}_j^T \mathbf{B}\|_\infty^2 \ \leq \ C_{X3} C_1^2 L^{-4}. \end{aligned}$$

Therefore, there is a positive constant C_2 such that

$$\|\overline{\beta}_j(T) - \overline{\boldsymbol{\gamma}}_j^T \mathbf{B}(T)\| \leq C_2 L^{-2}.$$

This implies that

$$\|\overline{\beta}_j(T)\| - C_2 L^{-2} \leq \left\{ \overline{\boldsymbol{\gamma}}_j^T \mathrm{E}\{\mathbf{B}(T)\mathbf{B}(T)^T\} \overline{\boldsymbol{\gamma}}_j \right\}^{1/2} \leq \|\overline{\beta}_j(T)\| + C_2 L^{-2}. \tag{34}$$

The desired result follows from (26) and (34). □

Proof of Lemma 3. Recall the notation defined at the beginning of this section. First we deal with $|d_{lS}|$ and $|d_{lSn} - d_{lS}|$.

We have $|d_{lS}| \leq C_1 (p_0/L)^{1/2}$ from the definition of the B-spline basis. As in the proof of Lemma 2 of [5], we have

$$|d_{lSn} - d_{lS}| \leq \delta (L p_0)^{1/2} / n$$

uniformly in $S \subsetneq S_0$ and $l \in S^c$ with probability

$$1 - C_2 p_0 L \exp\{-\delta^2 (C_3 n L^{-1} + C_4 \delta)^{-1}\} \times p \exp(p_0 \log 2),$$

where C_2, C_3, and C_4 are positive constants independent of p_0, L, n, p, and δ.

By combining the above results, (29), and Lemma 1, we obtain

$$|D_{lSn}^{-1}(d_{lSn} - d_{lS})| \leq C_5 L^{3/2} p_0^{1/2} \delta / n \tag{35}$$

and

$$|(D_{lSn}^{-1} - D_{lS}^{-1}) d_{lS}| \leq |D_{lS}^{-1}||D_{lSn} - D_{lS}||D_{lSn}^{-1}||d_{lS}| \leq C_5 L^{3/2} p_0^{3/2} \delta / n \tag{36}$$

uniformly in $S \subsetneq S_0$ and $l \in S^c$, with probability given in the lemma. Note that C_5 is independent of p_0, L, n, and δ. Hence the desired result follows from (35) and (36). □

References

[1] A. Antoniadis, I. Gijbels and A .Verhasselt, Variable selection in varying-coefficient models using P-splines, Journal of Computational and Graphical Statistics 21 (2012) 638-661.

[2] P. J. Bickel, Y. A. Ritov and A. B. Tsybakov, Simultaneous analysis of Lasso and Dantzig selector, Annals of Statistics 37 (2009) 1705-1732.

[3] E. Candes and T. Tao, The Dantzig selector: Statistical estimation when p is much larger than n, Annals of Statistics 35 (2007) 2313-2351.

[4] J. Chen and Z. Chen, Extended Bayesian information criteria for model selection with large model spaces, Biometrika 95 (2008) 759-771.

[5] M.-Y. Cheng, T. Honda, J. Li and H. Peng, Nonparametric independence screening and structural identification for ultra-high dimensional longitudinal data, forthcoming in Annals of Statistics and at arXiv preprint arXiv:1308.3942 (2014).

[6] B. Efron, T. Hastie, I. Johnstone and R. Tibshirani, Least angle regression (with discussions), Annals of Statistics 32 (2004) 407-499.

[7] J. Fan, Y. Feng and R. Song, Nonparametric independence screening in sparse ultra-high-dimensional additive models, Journal of the American Statistical Association 106 (2011) 544-557.

[8] J. Fan and R. Li, Variable selection via nonconcave penalized likelihood and its oracle properties, Journal of the American Statistical Association 96 (2001) 1348-1360.

[9] J. Fan and J. Lv, Sure independence screening for ultrahigh dimensional feature space, Journal of the Royal Statistical Society: Series B 70 (2008) 849-911.

[10] J. Fan, Y. Ma and W. Dai, Nonparametric independence screening in sparse ultra-high dimensional varying coefficient models, forthcoming in Journal of the American Statistical Association.

[11] J. Fan and R. Song, Sure independence screening in generalized linear models with NP-dimensionality, Annals of Statistics 38 (2010) 3567-3604.

[12] J. Fan, L. Xue and H. Zou, (2014). Strong oracle optimality of folded concave penalized estimation, Annals of Statistics 42 (2014) 819-849.

[13] J. Fan and W. Zhang, Statistical methods with varying coefficient models, Statistics and its Interface, 1 (2008), 179-195.

[14] D. Harrison and D. Rubinfeld, Hedonic housing prices and the demand for clean air, Journal of Environmental Econometrics and Management 5, 81-102.

[15] J. Z. Huang, C. O. Wu and L. Zhou, Polynomial spline estimation and inference for varying coefficient models with longitudinal data, Statistica Sinica 14 (2004) 763-788.

[16] W. Y. Hwang, H. H. Zhang and S. Ghosal, FIRST: Combining forward iterative selection and shrinkage in high dimensional sparse linear regression, Statistics and Interface 2 (2009) 341-348.

[17] H. Lian, Variable selection for high-dimensional generalized varying-coefficient models, Statistica Sinica 22 (2012) 1563-1588.

[18] H. Lian, Semiparametric Bayesian information criterion for model selection in ultra-high dimensional additive models, Journal of Multivariate Analysis 123 (2014) 304-310.

[19] J. Liu, R. Li, R. Wu, Feature selection for varying coefficient models with ultrahigh dimensional covariates, Journal of the American Statistical Association 109 (2014) 266-274.

[20] S. Luo and Z. Chen, Sequential Lasso cum EBIC for feature selection with ultra-high dimensional feature space, forthcoming in Journal of the American Statistical Association.

[21] L. Meier, S. van de Geer and P. Bühlmann, The group lasso for logistic regression, Journal of the Royal Statistical Society: Series B 70 (2008) 53-71.

[22] H. S. Noh and B. U. Park, Sparse variable coefficient models for longitudinal data, Statistica Sinica 20 (2010) 1183-1202.

[23] P. Radchenko and G. M. James, Improved variable selection with forward-lasso adaptive shrinkage, Annals of Applied Statistics 5 (2011) 427-448.

[24] L. L. Schumaker, Spline Functions: Basic Theory 3rd ed, Cambridge University Press, Cambridge, 2007.

[25] R. Song, F. Yi and H. Zou, On varying-coefficient independence screening for high-dimensional varying-coefficient models, Statistica Sinica 24 (2014) 1735-1752.

[26] Y. Tang, H. J. Wang, Z. Zhu, X. Song, A unified variable selection approach for varying coefficient models, Statistica Sinica 22 (2012) 601-628.

[27] R. J. Tibshirani, Regression shrinkage and selection via the Lasso, Journal of the Royal Statistical Society: Series B 58 (1996) 267-288.

[28] L. Wang, H. Li and J. Z. Huang, Variable selection in nonparametric varying-coefficient models for analysis of repeated measurements, Journal of the American Statistical Association 103 (2008) 172-183.

[29] F. Wei, J. Huang and H. Li, Variable selection and estimation in high-dimensional varying-coefficient models, Statistica Sinica 21 (2011) 1515-1540.

[30] L. Xue and A. Qu, Variable selection in high-dimensional varying-coefficient models with global optimality, Journal of Machine Learning Research 13 (2012) 1973-1998.

[31] Y. Xia, W. Zhang and H. Tong, Efficient estimation for semivarying-coefficient models, Biometrika 91 (2004) 661-681

[32] M. Yuan and Y. Lin, Model selection and estimation in regression with grouped variables, Journal of the Royal Statistical Society: Series B 68 (2006) 49-67.

[33] C. H. Zhang, (2010). Nearly unbiased variable selection under minimax concave penalty, Annals of Statistics 38 (2010) 894-942.

[34] W. Zhang, S. Lee and X. Song, Local polynomial fitting in semivarying coefficient model, Journal of Multivariate Analysis 82 (2002) 166-188.

[35] P. Zhao and L. Xue, Variable selection in semiparametric regression analysis for longitudinal data, Annals of the Institute of Statistical Mathematics 64 (2012) 213-231.

[36] H. Zou, The adaptive Lasso and its oracle properties, Journal of the American Statistical Association 101 (2006) 1418-1429.